arXiv:1611.08154v1 [cs.HC] 24 Nov 2016

AutoGain: Adapting Gain Functions by Optimizing Submovement Efficiency

Byungjoo Lee[1,2]
byungjoo.lee@kaist.ac.kr

Mathieu Nancel[1]
mathieu@nancel.net

Antti Oulasvirta[1]
antti.oulasvirta@aalto.fi

[1]Department of Communications and Networking, Aalto University, Finland
[2]Graduate School of Culture Technology, KAIST, Republic of Korea

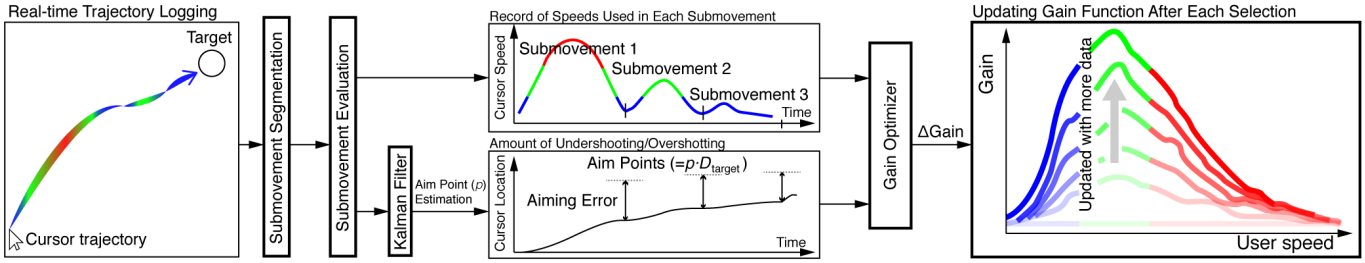

Figure 1. AutoGain improves performance on indirect pointing devices unobtrusively by using a novel submovement-level tracking+optimization approach. It gradually updates the control-to-display (CD) gain function based on analysis of speed and error in a user's submovements. It updates the CD gain function trying to minimize expected aiming error for typical submovements.

ABSTRACT

A well-designed control-to-display (CD) gain function can improve pointing performance with an indirect pointing device such as a trackpad. However, the design of gain functions has been challenging and mostly based on trial and error. *AutoGain* is an unobtrusive method to obtain a gain function for an indirect pointing device in contexts where cursor trajectories can be tracked. It gradually improves pointing efficiency by using a novel submovement-level tracking+optimization technique. In a study, we show that AutoGain can produce gain functions with performance comparable to commercial designs in less than a half hour of active use. This is attributable to reductions in aiming error (undershooting/overshooting) for each submovement. Our second study shows that AutoGain can be used on emerging input devices (here, a Leap Motion controller) that do not have reference gain functions yet. Finally, we discuss deployment in a real interactive system.

Author Keywords

Pointing; Submovement; CD gain functions; Pointer acceleration; Human Performance; Pointing facilitation.

INTRODUCTION

This paper investigates a novel approach to improving control-to-display (CD) gain functions, a class of pointing facilitation techniques used in the vast majority of indirect input devices like mice and trackpads. These functions are used to determine the speed of a control point in display space (e.g. a cursor) from the user's motion speed in control space (e.g. the speed of a mouse). A gain function computes a scalar factor (or "gain") from the instantaneous input speed (v_{in}), that is multiplied to that input speed to obtain the speed of the control point (v_{out}): $v_{out} = f_{CD}(v_{in}) \times v_{in}$. The design of a gain function greatly affects the kinematics [16, 17] and performance [29, 4, 2, 22] of pointing actions. Improvements in selection time of up to 24 % were observed with speed-dependent gain functions [2], compared to constant gains.

Despite its importance, the design of gain functions has been based on either extensive trial and error [2] or iteration with heuristics [22, 30] (see Related Work). No automated or computer-assisted method exists for determining an efficient function for a given device (see [22]). The challenge is that the design space of gain functions is very large: in principle any continuous function can be used. However, to be efficient, a function must strike a balance between high and low gains to account for both speed and accuracy.

This paper presents *AutoGain*, a novel method to design gain functions that addresses these challenges (Figure 1). It iteratively adapts the CD gain function after each target acquisition, using the estimated accuracy of its submovements. These changes are designed to minimize aiming error (overshooting/undershooting), which is calculated with respect to the user's inferred aim point for each submovement within the last target acquisition (see the two graphs in the center of Figure 1). The gain function is treated as a discrete function, with input speeds "binned" like in a histogram, and the gain of each bin is modified based on the speeds used in the previous pointing task.

AutoGain can produce a function in a matter of thirty minutes of active use. Critically, unlike previous methods, it requires no human supervision and minimal initialization: after setting the overall speed of convergence, it starts from a constant function and adapts it based on automated observations and principles. This makes AutoGain easily deployable, and suitable for input devices for which no reference functions exist yet. The method can also adapt the gain function to abrupt changes in conditions, e.g. a new user, or a different device.

This approach is somewhat analogous to playing golf. The golfer (user) wants to relocate the ball (cursor) as close as possible to the hole (target) with as few hits (submovements) as possible. In this metaphor, AutoGain is like a caddy who constantly analyzes the player's previous shots to prepare a better set of clubs (gain function), that will reduce his aiming error in the future for similar holes (pointing difficulty).

We detail the AutoGain method, and report the results of two controlled studies. In the first one, AutoGain produced a gain function for a laptop trackpad, with performance close to that of a widely used commercial function (Mac OS X). In the second study, AutoGain produced gain functions for a less conventional input device (a Leap Motion) with performance comparable to previous work using hand-tuned gain functions. In both studies, convergence occurred in less than thirty minutes of use. We conclude by discussing how AutoGain could be integrated into real interactive systems.

RELATED WORK

Although virtually every study using an indirect pointing device has used some CD gain function, their design and adaptation is a relatively little studied topic (see [22, 30]).

Casiez and colleagues [4] presented a procedure to choose constant gain values. They define a range of gains between two endpoints: a minimum gain CD_{min} that allows the user to acquire the most distant targets without clutching, and a maximum gain CD_{max} that ensures that each pixel can be reached. Gains within this interval support both objectives.

Casiez and Roussel [2] reverse-engineered gain functions in existing operating systems. They found that all commercially deployed functions share some features, namely monotonous increase in the beginning and comparable maximum outputs. They also found differences, especially in minima, continuity, and shape of the functions. Empirical comparison among the functions showed that speed-dependent gain functions can be up to 24% faster than a constant CD gain.

Nancel and colleagues [21, 22] proposed a generic gain function based on the generalized logistic curve. It expresses four features of gain functions: the asymptotic minimum and maximum output gains, the abscissa (input velocity) of the curve's inflexion point, and its slope at that inflexion point. This function was successfully applied to translation- and rotation-based input channels, and later on to different input devices and interactive environments [11, 18]. However, tuning the parameters of the function is ad hoc: initial values are based on heuristics derived from [4], and suitability is left to the judgment of the designer. Earlier approaches combined simple sine-based CD gain functions with absolute position control [9, 14, 10], with an implicit velocity-based transition.

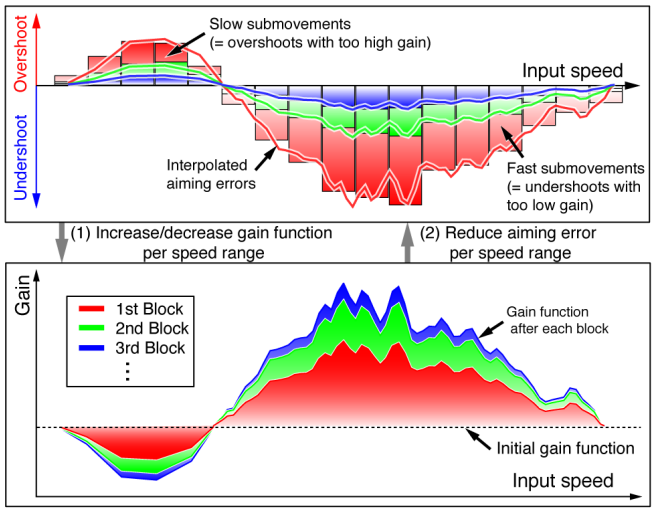

Figure 2. AutoGain improves a discrete gain function by updating gains locally per speed range. It segments a movement into submovements to update a profile of aiming errors, which it tries to reduce by adapting the gain response. Here, it is shown how AutoGain fixes too-low gains at high input speeds, and too-high gains at low input speeds.

Blanch and colleagues [1] proposed an unobtrusive method to change the CD gain during the target acquisition, based on the semantic importance of nearby targets. Their result showed that the visual information of target alone is sufficient for a user to select the target, even when the CD gain is changing based on target proximity and salience. We suppose that users are dealing with the unnaturalness by maintaining only lightweight representation on their reaching movements. This idea shed lights on possibility of submovement-level adaptation techniques.

In contrast to previous studies considering how the magnitude of input movement is being transferred, several studies have investigated other aspects of control-to-display relationships, such as angular deviation [29], coordinate disturbance [17], and movement direction [16]. Results include accurate and faster drawing [17], reduced overshoot [16], and better experience for motor impaired users [29].

To sum up, it has been shown that speed-dependent gain functions are superior to constant functions. However, the design of gain functions has been limited to heuristic–iterative approaches. Research shows that there is value in analyzing pointing kinematics in their design.

OVERVIEW OF TECHNIQUE

AutoGain is built under the core assumption that *the aiming error of individual pointing submovements is an indication of the adequacy of the gain function*: if a submovement ends before its intended aim point (undershoot), then the gain function was too low and should be increased; if it ends further than its intended aim point (overshoot), then the gain function was too high and should be decreased. In both cases, the change is proportional to the amount of aiming error. Auto-

Gain applies this principle in a speed-dependent manner: it treats gain functions as series of gains associated to discrete intervals of input speeds, and only alters a gain when the corresponding "speed bin" was used in a pointing task. This principle is summarized in Figure 2.

AutoGain subscribes to a local optimization scheme that assumes that there exists a case-specific optimum gain function (or several optimal gain functions on a Pareto front) for a user performing pointing tasks of certain difficulty and scale ranges, with a given pointing device. It also assumes that this function is reachable from an initial constant gain function, by means of a sequence of local updates to the function. Moreover, when updates are based on repeated observations of the user, a reasonable estimate of the optimum function can be obtained despite several sources of variability in the human motor system.

To estimate the optimum, AutoGain builds on four assumptions in earlier theories of human motor control:

Submovement decomposition: An aimed movement can be divided into one or more submovements [27, 6, 20] using local accelerations and decelerations in the speed profile.

Implicit aim point: A submovement has an implicit aim point located at a fraction (p) of the remaining distance to the target (D_{target}) at its beginning [6]:

$$D_{aim} = p \cdot D_{target} \tag{1}$$

With respect to the aim point, the aiming error R of a submovement is defined as:

$$R = (D_{aim} - D_c) \tag{2}$$

With D_c the distance moved during the submovement (projected, see Figure 4). Negative errors indicate overshooting, positive errors indicate undershooting.

Stochastic noise: Due to stochastic noise in the motor system, submovement endpoints are distributed around the center of the aim point [20].

Intermittent control: The next submovement starts from where the previous was left [6, 20].

These assumptions extends tenets of two well-known models of aimed movements. In Crossman and Goodeve's *deterministic iterative corrections model* [6], aimed movements are modeled as series of ballistic, open-loop submovements aimed at a constant fraction (p) of the remaining distance to the target. However, because empirical data showed large variations in the duration and aim point of these submovements, an extension of the idea was proposed. Meyer's *stochastic optimized submovement model* [20] assumes that neuromotor noise causes the primary submovement to either undershoot or overshoot the target, which requires a corrective submovement to finally reach the target center. It assumes intermittent feedback control, where each submovement is based on the error at the end of the previous one.

AutoGain's objective function is to minimize aiming errors. To this end, it applies an iterative search method similar to a gradient descent local optimization scheme. It changes (decrease/increase) the gain function proportionally to the amplitude of aiming error (overshooting/undershooting), and only for the specific speeds actually used in the submovement.

IMPLEMENTATION

AutoGain updates the gain function after every target selection. The updating procedure consists in five steps (see Figure 1). First, it records a pointing trajectory from onset until the target selection. Second, it segments the trajectory into submovements using kinematic criteria. Third, it filters submovements out based on their trajectory and dynamic properties. Fourth, it computes the amount of aiming error (undershooting/overshooting) for each submovement, relative to its estimated aim point, as well as the input speeds that were used in that submovement. Fifth, it updates the gain function around these input speeds, using the amplitude and nature of these errors, following a local optimization scheme.

Step 1: Real-Time Trajectory Logging

AutoGain logs two different time vectors in real-time, for each input event t: (1) the raw input stream (dx_t, dy_t) (in counts) from the input device, and (2) the cursor trajectory $(x_{c,t}, y_{c,t})$ (in pixels). The raw input stream is used to obtain records of movement speeds in submovements. The cursor trajectory is used to segment submovements. We express the relationship between inputs and outputs as follows:

$$\begin{aligned} v_t &= C_{in}\sqrt{dx_t^2 + dy_t^2} \\ x_{c,t+1} &= x_{c,t} + C_{out} \cdot (C_{in} \cdot dx_t \cdot G[v_t]) \\ y_{c,t+1} &= y_{c,t} + C_{out} \cdot (C_{in} \cdot dy_t \cdot G[v_t]) \end{aligned} \tag{3}$$

$$\text{with } C_{in} = \text{Freq}_{in}/\text{Res}_{in} \text{ and } C_{out} = \text{Freq}_{in}/\text{Res}_{out}$$

C_{in} is a factor converting unit of raw input (count) to m/s, which is determined from the input device's resolution (Res_{in}, here in points per millimeters) and update frequency (Freq_{in}, here in number of events per millisecond) [2]. C_{out} uses the opposite principle to convert transformed input movements in m/s into pixels translations, using the same update frequency Freq_{in} and the resolution of the display (Res_{out}, in pixels per millimeter). AutoGain treats gain functions as arrays of gains associated to discretized intervals of input speed. Gains values ($G[v_t]$) are therefore interpolated when v_t is not one of the discretized input values.

Step 2: Submovement Segmentation

From the beginning of a movement to the selection of the target, the coordinates are segmented into submovements based on local extrema in the cursor speed profile [7], using the `Persistence1D` [15] algorithm. It first smooths the speed profile (red curve in Figure 3) using a Gaussian kernel filter ($\sigma = 3$), then returns all pairs of minima and maxima that exceed a pre-defined persistence value (0.2). In most cases, resolution of raw inputs are available in integer steps so we set value of persistence less than 1.0 to ensure enough sensitivity on smaller submovements even after the smoothing process.

After identifying the local minima and maxima in the speed profile, each neighboring minimum-maximum-minimum

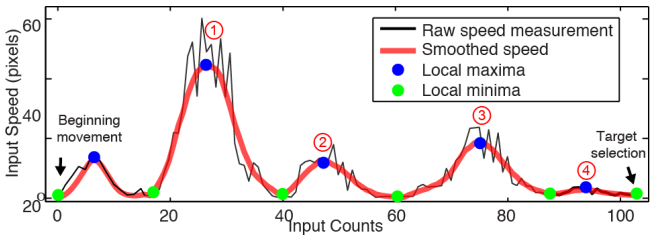

Figure 3. After smoothing the input speed curve of a trial movement, `Persistence1D` scans every minimum-maximum-minimum triplet to segment submovements. Submovements preceding the highest maximum (here ①) are excluded.

triplet is considered to be a possible submovement. As in [7], we only consider submovements from the highest local maximum (① in Figure 3), which is assumed to be the initial ballistic movement to the target. This helps exclude possible non-aiming movements during the trial.

This process returns the cursor coordinates corresponding to each of the identified submovements, i.e. that occurred between the timestamps of the corresponding speed minima.

Step 3: Submovement Evaluation

AutoGain then identifies some unwanted characteristics of the submovements' trajectories: *unaimed*, *interrupted*, and *non-ballistic*. The goal of this classification is to filter out submovements that are likely to introduce noise in the updates of the gain function, or of the aim point (see Table 1).

We define *unaimed* submovements using two tajectory properties: (1) the maximum angular deviation, and (2) the amount of overshoot (see Figure 4). Maximum angular deviation is defined as the maximum angle between the line joining the first and last points of the submovement's trajectory, and the line joining the first and any other point of the submovement's trajectory. The amount of overshoot is defined as:

$$\max(D_c - D_{target}, 0) \tag{4}$$

Any submovement that satisfies at least one of the following conditions is marked as *unaimed*:

- Maximum angular deviation $> 45°$.
- Amount of overshoot $> 0.5 \times D_{target}$.

Submovements are considered *interrupted* when they fall distinctly short (less than halfway) from the remaining distance to the target, or when they include clutching. Clutching can be determined with more or less certainty depending on the input device. Some devices allow straightforward detection

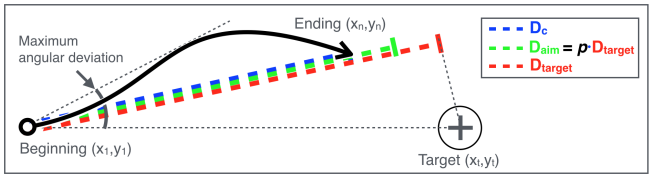

Figure 4. AutoGain assumes that each submovement has an implicit aim point (D_{aim}) located at a fraction (p) to the remaining target distance (D_{target}). In this figure, a submovement is undershooting, which is marked with a positive aiming error $R = (D_{aim} - D_c)$.

Type of Submovement	Ballistic ($D_{aim} = pD_{target}$)	Non-ballistic ($D_{aim} = D_{target}$)
Normal	Update gain: Yes Update aim point: Yes	Update gain: Yes Update aim point: No
Interrupted	Update gain: Yes Update aim point: No	Update gain: Yes Update aim point: No
Unaimed	Update gain: No Update aim point: No	Update gain: No Update aim point: No

Table 1. Depending on the type of submovements, details of gain and aim point updates changes. The general objective of this classification is to exclude submovements with insufficient quality (noisy and unaimed movements) in order to facilitate the convergence of AutoGain.

of clutching by providing touch-down and touch-up events (e.g., styluses). Otherwise, predefined temporal thresholds can be used to detect the resetting of the end-effector, by measuring the time between two consecutive sensor events. In our implementation with a trackpad, we categorize a submovement as clutching when any interval of sensor events exceeds a predefined temporal threshold (130 ms). The last submovement of a task was never considered as clutching. For simplicity, submovements that are neither *unaimed* nor *interrupted* are deemed "normal".

Submovements that happen after the second normal submovement are assumed to be *non-ballistic*, i.e. continuously controlled. This assumption is based on the optimal submovement theory [20], which states that in ideal cases two submovements are enough to reach to a target. Although there are still debates on the ballistic nature of the secondary submovement[8, 27], in this work we make the assumption that the secondary submovement is ballistic.

That classification is used to decide whether AutoGain will include a submovement in future updates of the gain function and the aim point, as reported in Table 1. *Unaimed* submovements are excluded from the updates due to their poor aiming quality. *Non-ballistic* submovements are considered aimed directly at the target ($D_{aim} = D_{target}$) rather than at the user's implicit aim point ($D_{aim} = p \cdot D_{target}$), and therefore do not update the estimated aim point. *Interrupted* submovements are treated as aiming for the estimated aim point ($D_{aim} = p \cdot D_{target}$), but falling short due to insufficient gain, and therefore do not update the estimated aim point.

Step 4: Speed Profiles and Aiming Errors

AutoGain measures the aiming error of interrupted and normal submovements. We defined in Equation 2 the aiming error R_i of a submovement i as the (projected) distance remaining to its estimated aim point at the end of the submovement (see Figure 4): $R_i = (D_{aim,i} - D_{c,i})$. In non-ballistic submovements, the aim point is assumed to be the center of the target. In ballistic movements, we assume that the user is aiming at a point located before the target, at a certain proportion p of the remaining distance. This proportion is likely to vary with the gain function and the user's strategy, so AutoGain approximates it from repeated observations using a Kalman filter [13]. The filter is updated after each *normal+ballistic* submovement.

$$p_i = f_{\text{Kalman}}\left(\frac{D_{target,i} - D_{c,i}}{D_{target,i}}, p_{i-1}\right) \tag{5}$$

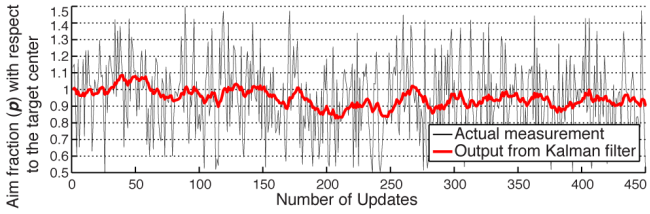

Figure 5. The aim proportion p is continuously updated by a Kalman filter applied after each ballistic submovement from actual measurements. The graph is showing the updates from participant 6 in Experiment 1 with AUTOGAIN condition (trial 1 to 60 in the first BLOCK).

This filtering is based on the assumption of stochastic noise in human movement, whose distribution is centered at the aim point. From the repeated observations, the filter removes stochastic noise similarly to a low pass filter (see Figure 5). This supposes possible changes in the implicit aim point is slow. After estimating the aim fraction p_i of the current submovement, the aiming error is calculated using Equation 2: $R_i = (p_i \cdot D_{target,i} - D_{c,i})$.

AutoGain also records the input speeds of every event in each submovement, in the form of a boolean array $S_i[V]$. The range of possible input speeds is binned into J speed intervals ($V_j = [v_j, v_{j+1}], j \in [0, J-1]$) of constant width $w = v_{j+1} - v_j$. The corresponding array entry is false unless at least one event in the corresponding submovement had an instantaneous speed v within that interval:

$$S_i[V_j] = \begin{cases} 1, & \text{if } \exists\, v : v_j < v < v_{j+1} \text{ in submovement } i \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

R_i and $S_i[V]$ are used in the gain optimization process.

Step 5: Gain Update

The principle behind AutoGain, after each trial, is to increase (respectively decrease) the gain function when submovements in that trial undershot (respectively overshot) their estimated aim points, for the input speeds that were used during these submovement. AutoGain considers a gain function as a series of gain values associated to binned intervals of speeds input, using the same partitioning as in $S_i[V]$ (Equation 6). The gain function for the next trial $G_{t+1}[V]$ is updated from the gain function of the current trial $G_t[V]$ as follows:

$$G_{t+1}[V] = G_t[V] + \sum_{i=1}^{N_t} \Delta_i[V] \tag{7}$$

With N_t the number of submovements in the trial t, and Δ_i the corrections calculated from each submovement's aiming error. The amplitude of gain change $\Delta_i[V_j]$ for each speed bin is calculated using the amount of aiming error R_i, multiplied by a constant C that defines the overall rate of gain change.

As stated earlier, only the speed bins that have been used in the submovements of the previous pointing task will be updated ($S_i[V]$ in Equation 9). However, a given speed interval can be used by more than one submovement within the same pointing task. To avoid over-favouring the speeds most commonly used, AutoGain prioritizes submovements in descending chronological order, on the principle that submovements that occurred earlier in a trial can be sufficiently represented by their faster speed components. The gain change of each speed bin V_j is therefore calculated in reverse order from the last submovement N_t, allowing only one change per bin (Equation 8). $\Delta_i[V]$ is calculated as expressed in Equation 9.

$$I_i[V] = \prod_{k=(i+1)}^{N_t} (1 - S_k[V]) \tag{8}$$

$$\Delta_i[V] = C \cdot R_i \cdot S_i[V] \cdot I_i[V] \tag{9}$$

Change Rate C

In effect, the AutoGain optimization process involves two interdependent active components: (1) gain optimization and (2) human skill acquisition. The system adapts the gain function to the user movements, then the user adapts his movements to the changes in the gain function, and so on. It is therefore crucial that the updates to the gain function occur fast enough to ensure a realistically short calibration process, but slow enough to allow the user to adapt to them.

This is implemented in AutoGain through the parameter C in Equation 9. If C is large, the user will not have enough time to follow up the updates in the gain function and the system will become soon unstable. However, if C is small, it will take too much time for the performance to converge. Overall, the effect of C on convergence speed will depend on the characteristics of the interactive system, like the scale difference between the input and output resolutions.

C expresses the relationship between the amount of gain change δ_g that should occur after M submovements of average aiming error μ_R, for a given speed bin:

$$C = \frac{\delta_g}{M \cdot \mu_R} \tag{10}$$

In future applications of AutoGain, we recommend to start with an initial C value of 5×10^{-5} (in mm^{-1}), then increase or decrease the value through pilot testing based on the time pressure imposed on the optimizer and user. With this initial value, AutoGain changes the gain by 1 every 10,000 submovements with an average aiming error of 2 mm, or every 1000 submovements with an average error of 2 cm.

STUDY 1: TRACKPAD GAIN FUNCTION

We ran a study to assess how AutoGain fares in a typical laptop + trackpad setup, how long it takes to converge, and how the resulting gain function compares to a well-established baseline function (Mac OS X) in shape and performance.

Participants

We recruited 11 paid participants (4 females) aged 22 to 38 (μ 27.4 years old, σ 5.3) from the local university. They were all regular MacBook trackpad users (for 4.5 years on average, σ 2.2) and used a trackpad 5 hours a day on average (σ 3.1). Only one participant was left-handed, but typically used her right hand to control the trackpad. We also asked their current trackpad setting in Mac OS X (one of 10 positions on a slider,

the higher the faster): 4 participants used the 4th position, 2 used the 6th, 4 used the 7th, and 1 used the 8th.

Design and Dependent Variables

The experiment followed a within-subject design with one independent variable: *gain condition* = {REFERENCE, AUTOGAIN}. In the REFERENCE condition, we replicated each participant's everyday pointer acceleration setting (see above) using the `libpointing` [2] library. In the AUTOGAIN condition, the gain function was initially set to a constant gain of 1 ($G[v_t] = 1\ \frac{m.s^{-1}}{m.s^{-1}}$). This function was however updated after each trial using the AutoGain method.

Participants completed 800 trials for each *gain condition*. We counterbalanced the ordering of conditions given to the participants. For performance analysis, we blocked those trials into 10 *block* conditions of 80 trials with randomized Fitts's Indexes of Difficulty (ID). We used three dependent variables to compare REFERENCE and AUTOGAIN: *trial completion time*, *error rates*, and subjective assessments using the NASA-TLX form [12].

Task and Procedure

The task consisted in selecting circular targets on a laptop screen using the embedded trackpad, in the two different *gain condition*s. Participants were instructed to perform the tasks as quickly and as accurately as possible. In the AUTOGAIN condition, participants were informed about the gain adaptation. They were also warned that pointing might feel slow or awkward in the beginning, and instructed to try to perform normally regardless. Participants were allowed use either tapping or the embedded physical button to select the targets.

The target was a red disk presented on a black background. To obtain a comprehensive picture of the observed effects, we randomized the target diameter ($[2, 11.5]$ mm), the orientation between two consecutive targets ($[0, 2\pi]$ rad), and the ID of each task ($[2, 5.5]$ bits). We used the following randomization process after each successful click (x, y):

1 : $ID \leftarrow$ random(2, 5.5);
2 : **do:**
3 : $W_c \leftarrow$ random(2, 11.5);
4 : $x_c \leftarrow$ random(0, screen.width);
5 : $y_c \leftarrow$ random(0, screen.height);
6 : $ID_c \leftarrow \log_2\left(1 + \frac{\text{dist}[(x,y),(x_c,y_c)]}{W_c}\right)$;
7 : **until** $|ID_c - ID| < 0.1$

This ensured uniform distributions of IDs and orientations (see Figure 6), providing comparable performance datasets. Note that this process favors small target sizes, that have more chances to fit in the screen when orientation is random.

Participants sat on a regular office chair that they could adjust, and used a laptop placed on a desk. They first filled a preliminary questionnaire about their regular MacBook trackpad usage. Every 80 selections, participants were instructed to take a break and to answer a NASA-TLX form about their performance since the last break. The experiment lasted about one hour per participant.

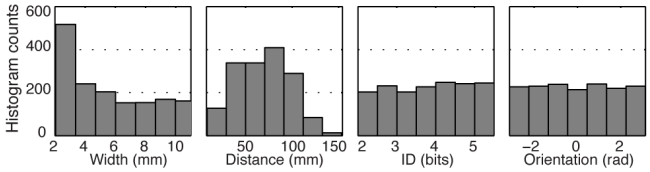

Figure 6. Study 1: We used a randomization scheme to obtain uniform distributions for ID and target orientation.

Apparatus

We ran the experiment on a MacBook Pro laptop (2012 version) running Mac OS X 10.11, using the integrated trackpad for input. The size of the display was 35.8 cm 24.7 cm (1280 $\times$ 800 pixels). The experiment was coded in C++. The optimization speed parameter C in Eq. (7-9) was set to 6.4×10^{-5} mm^{-1} after pilot tests, to ensure that the transitions between two instances of the gain function remain mostly unnoticed but ensure short duration of convergence. The REFERENCE functions were obtained from `libpointing` [2], which was also used for cursor coordinate calculations in this *gain condition*. We used `libpointing`'s 'subpixel' option that transfers the remainder of the last calculated (floating) cursor coordinates to the next time step. We implemented the same subpixel mechanism in AutoGain, and discretized the range of input speeds into bins of 0.0079 m/s as REFERENCE condition. We set the processing noise parameter of the Kalman filter at 0.2, and the sensor noise parameter at 40.0. The refresh rate of the display was 60 fps.

Results

Dataset

We excluded one participant from our dataset, who displayed irregular initial performance in both *gain condition*s as well as inconsistent aiming behavior (high variance in the p parameter, see Equation 1). She also performed the study using her right hand despite describing herself left-handed overall. Perhaps as a consequence of using her non-dominant hand, her ballistic movements were notably slower. Her average input speed for the first BLOCK was only 15% higher in AUTOGAIN than in REFERENCE, while the other participants were 96% faster with AUTOGAIN in that block. As a result, while AutoGain did improve her performance over time, that improvement was markedly slower. We exclude that participant from the following analyses (N=10), and take note that the consistency of movement can play an important role in the improvement rate of AutoGain.

We analyzed 8,000 trials per *gain condition*, corresponding to 19,766 submovements in the REFERENCE condition (2.47 per trial) and 19,828 submovements in the AUTOGAIN condition (2.48 per trial). In the REFERENCE condition, 4349 submovements (22.0% of all submovements, 0.54 times per trial on average) were evaluated as *interrupted*, and 4774 submovements (24.1% of all submovements, 0.60 times per trial) were evaluated as *unaimed* submovements, and 6031 submovements (30.5% of all submovements, 0.75 times per trial) were evaluated as *non-ballistic*. In the AUTOGAIN condition, 4275 submovements (21.6%, 0.53 times per trial) were *interrupted* submovements, and 3577 (18.0 %) were evaluated

as *unaimed* submovements, and 6769 submovements (34.1% of all submovements, 0.75 times per trial) were evaluated as *non-ballistic*. The overall average of aim point (p) estimated from Kalman filter was 0.94 (σ 0.033) for AUTOGAIN condition. In what follows, the completion time and error rates are averaged by blocks of 80 trials before analysis.

In the REFERENCE condition, the participant's typical trackpad setting (which was replicated in that condition) had no significant effect on trial completion time (F(1,8)=0.770, p=0.406). The ordering of gain conditions also had no significant effect on overall trial completion time (F(1,8)=0.187, p=0.677). Also the ordering did not affect the trial completion time of each gain condition (F(1,8)=0.009, p=0.928) or each block (F(1,8)=1.750, p=0.222) separately. So we exclude above factors in the following analysis.

Error Rates

The mean error rate was 9.3% (σ 2.2) for REFERENCE and 9.3% (σ 2.8) for AUTOGAIN. We attribute that overall high error rate to the higher number of small targets generated by the target randomization scheme (Figure 6), as has been reported in previous studies [28, 31]. We found no significant effect of *gain condition* (F(1,9)=0.001, p=0.984) or *block* (F(1,9)=1.261, p=0.29) on error rate. We also found no interaction effect between *block* and *gain condition* on error rate (F(9,81)=1.02, p=0.432). We conclude that the participants maintained comparable accuracy throughout the blocks.

Trial Completion Time

We found a significant effect of of *gain condition* on trial completion time (F(1,9)=40.68, p<0.001). The overall average completion time was 813.5 ms (σ 30.7) for REFERENCE and 869.1 ms (σ 26.4) for AUTOGAIN. This result is straightforward: in the AUTOGAIN condition, participants started from a slow gain of 1 $\frac{m.s^{-1}}{m.s^{-1}}$. However, the overall difference was only about 50 ms. This is on a similar level with a difference reported between OS X function and Windows in a previous study ($= 50$ ms)[2].

The effect of interaction between *block* and *gain condition* on trial completion time was significant (F(9,81)=10.24, p<0.001). Pairwise comparison showed that the differences in trial completion time between *gain condition*s are becoming not significant from BLOCK 7 (p=0.35, see Figure 7). Completion times per block were as follows:

Block #	REFERENCE (ms)	AUTOGAIN (ms)	p
7	799.35 (σ 34.8)	816.95 (σ 26.4)	.35
8	806.44 (σ 28.8)	835.34 (σ 27.8)	.065
9	814.28 (σ 38.9)	833.54 (σ 32.4)	.141
10	799.32 (σ 30.2)	841.31 (σ 32.4)	.137

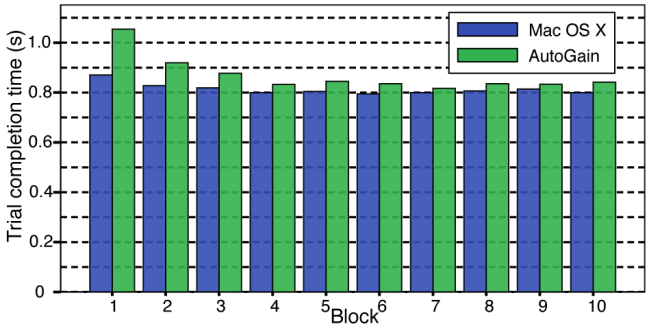

Figure 7. Study 1: The average difference in trial completion time between AUTOGAIN and REFERENCE was 56 ms. Moreover, the difference became insignificant from BLOCK 7 (p=0.35) per block.

Workload Metrics

For simplicity we report Raw TLX values [12]. We found a significant main effect of *gain condition*s on mental demand (F(1,9)=6.40, p=0.032), physical demand (F(1,9)=21.69, p=0.001), and effort (F(1,9)=9.30, p=0.014). However, from the interaction effect between *block* and *gain condition*, these differences pertain to the beginning half of the study and become insignificant throughout the study (see Table 2). Differences were not significant for temporal demand (F(1,9)=2.97, p=0.12), performance (F(1,9)=1.04, p=0.37), and frustration (F(1,9)=4.47, p=0.064). The overall trend was plotted in Figure 8.

Measure	REFERENCE	AUTOGAIN	Insignificant from
Mental demand	4.7 (σ 0.9)	6.4 (σ 1.3)	BLOCK 6 (p=0.062)
Physical demand	5.2 (σ 0.8)	7.9 (σ 1.2)	BLOCK 7 (p=0.063)
Effort	7.4 (σ 1.3)	9.0 (σ 1.4)	BLOCK 6 (p=0.133)

Table 2. Study 1: mean and standard deviation for overall TLX measure in both *gain condition*s. In every measure, the difference becomes insignificant after a certain BLOCK.

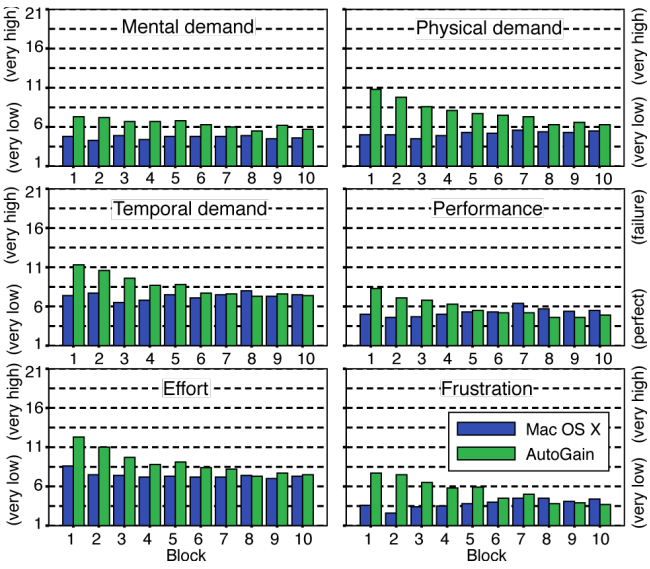

Figure 8. Study 1: The difference in subjective ratings between REFERENCE and AUTOGAIN becomes insignificant after 6 to 7 blocks in the trial. Users reported performance being higher in AUTOGAIN than in REFERENCE after BLOCK 7.

Discussion

The driving idea behind AutoGain is to reduce aiming errors as a means to improve pointing performance. Figure 9 shows that these errors are indeed being reduced over time. Performance between AutoGain and the baseline became similar at the same time as aiming errors, around BLOCK 7.

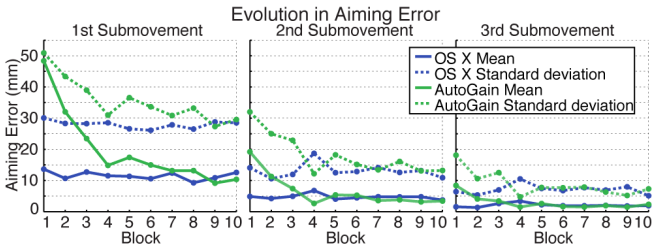

Figure 9. In Study 1, AutoGain gradually reduced aiming error in submovements to a similar level with the baseline (Mac OS X in-built gain function). Some fluctuating dynamics in aiming error were observed in BLOCK 5 and BLOCK 8 (black arrows).

Figure 10 shows the gain functions obtained by AUTOGAIN alongside the Mac OS X gain functions used in REFERENCE. Empirical measurements show that AutoGain can produce a CD gain function with performance comparable to commercial solutions within 30 minutes, both in terms of speed and accuracy. Participants also reported similar workload ratings by the end of the study. Figure 10 (Bottom) illustrates how the input speeds used in AUTOGAIN and REFERENCE became similar by the end of the experiment, especially between 0 and 0.3 m/s. This range also corresponds to the regions of the gain functions where AUTOGAIN is the most similar to the REFERENCE, both in value and slope.

The much steeper downward slope of the AutoGain-produced function after 0.3 m/s reflects the smaller number of input events featuring these input speeds: AutoGain updates the gain corresponding to an input speed interval only when this interval was used in the previous pointing task. However, the comparable completion times obtained between the two *gain conditions* tend to indicate that the gains above commonly used speeds is of lesser importance for overall performance.

STUDY 2: GAIN FUNCTION FOR AN EMERGING DEVICE

The objective of this study is to assess AutoGain's ability to produce usable gain functions on input devices that were not primarily designed for cursor control—in this case, a Leap Motion controller. This experiment design is similar to Study 1, with only the AUTOGAIN condition (no REFERENCE). The cursor was controlled using mid-air hand movements tracked by a Leap Motion controller, (software version 2.3.1+31549). We are mainly interested to see if AutoGain stabilizes to a gain function with satisfying performance and user feedback.

Participants

Six paid participants (4 females) were recruited, aged 22 to 30 years old (μ 25.5 years, σ 2.9). One participant was left-handed. One participant had 1 year of experience in mid-air pointing, mostly from gaming experience at home with another Leap Motion device.

Design

The experiment followed a within-subject design with one gain condition: AUTOGAIN, which is similar than in the previous study. The initial gain was set as a flat function with one-to-one mapping ($G[v] = 1\ \frac{m.s^{-1}}{m.s^{-1}}$) between the speed of

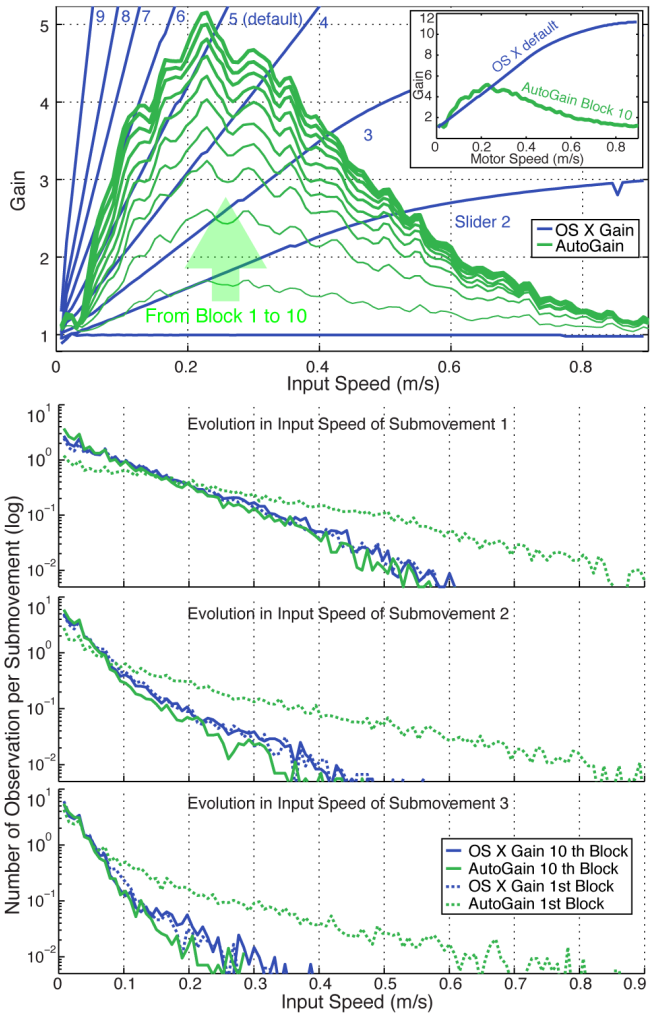

Figure 10. Top: AutoGain gradually updates the gain function until convergence (no more improvements in aiming error). Bottom: the evolution in user's input speed. After 10 blocks of trials, AutoGain participants were utilizing the same speed intervals as with the OS X function.

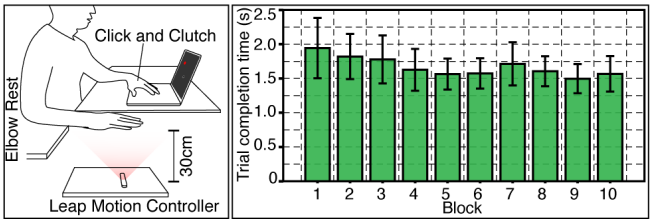

Figure 11. Apparatus and trial completion times in Study 2.

index finger and the cursor. Each participants was to perform 800 trials, which we blocked into 10 BLOCKS of 80 trials.

Task and procedure

The task (on screen) was the same as in Study 1. Participants used their dominant hand over a Leap Motion tracker to move the cursor, and their non-dominant hand over the built-in keyboard to click and clutch. Clicking was performed by pressing the space bar. The cursor only moved when the C key was held pressed, following the movements of the dominant index

finger; clutching was performed by releasing that key, relocating the finger, and resuming pointing by pressing C again. The finger movements were interpreted as parallel to the display, i.e. on a vertical plane. This ensured forward/up visual compatibility [24] between hand and cursor movements.

For each participant, we first calibrated the Leap Motion tracker using its inbuilt calibration software before starting the experiment. A research assistant then briefly demonstrated the pointing mechanism. Participants sat on a regular office chair that they could adjust, and used a laptop placed on a desk. They first filled a preliminary questionnaire asking about their previous experience with mid-air pointing.

During the tasks, participants rested their dominant elbow on a stand aside the desk (at desk height), with the Leap Motion device facing up below their hand and 30 cm below desk height (see Figure 11-left). In effect, their dominant hand rested 30 cm above the Leap Motion tracker, and their elbow was supported to reduce fatigue.

Participants were instructed to take 10 minutes of practice session with an initial constant gain of 1 before the main part of the experiment started. The gain function was not updated during the practice. During the main session, participants were invited to take breaks every 80 selection tasks. The experiment lasted about one hour per participant.

Apparatus

The experiment was coded in C++ and run on the same computer as in Study 1. We used the Leap Motion in "Robust tracking mode" and smoothened its raw input using the `1 € filter` [3]. Following Casiez *e*t al.'s tuning guidelines, we used 10^{-5} for minimum cutoff frequency and 0.05 for beta parameter.

The change rate parameter C was set to $3.6{\times}10^{-5}$ mm^{-1} after pilot tests; this is lower than Study 1 to account for the possible harder learning of the input method, compared to traditional trackpad input. We again set the processing noise parameter of the Kalman filter at 0.2, and the sensor noise parameter as 40.0. The refresh rate of the display was 60 fps.

For dynamic tracking of moving hands, the spatial resolution of Leap Motion controller was reported at around 0.7 mm [26]. Therefore, to avoid being sensitive to noise, we discretized the range of input speeds into bins of 1 mm/count (0.06 m/s considering Leap Motion's 60Hz frequency).

Results

Dataset

In total, 19,271 submovements (4.01 per trial on average) were observed, with 2,219 *interrupted* submovements (11.51% of all submovements, 0.46 times per trial on average), 7,832 *unaimed* submovements (40.64% of all submovements, 1.63 times per trial on average), and *non-ballistic* 9,142 submovements (47.44% of all submovements, 1.9 times per trial on average) according to our criteria. Note that we detected interrupted submovements in the same way as trackpad study. The overall average of aim point (p) estimated from Kalman filter was 0.94 (σ 0.038).

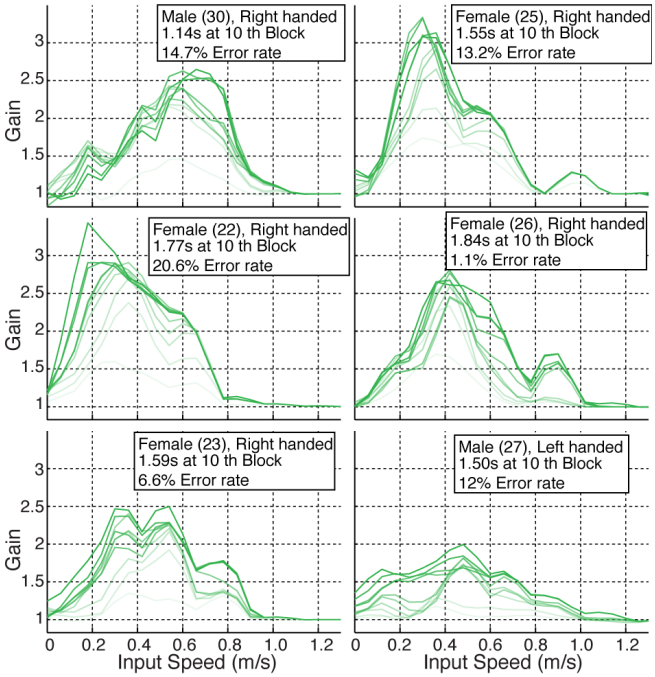

Figure 12. Study 2: AutoGain obtained gain functions with slightly different shapes for different participants.

Trial Completion Time and Error Rates

Trial completion time was 1.67 second on average (σ 0.29). Completion time varied from 1.94 second (σ 0.42) in BLOCK 1 to 1.57 second (σ 0.25) in BLOCK 10, for an improvement of 19.1%. The effect of BLOCK on trial completion time was significant ($F(9,45)=4.383$, $p<0.001$). Helmert contrast revealed that the improvement in trial completion time was close to being statistically significant for BLOCK 1 ($p=0.057$) and BLOCK 2 ($p=0.052$). Also for BLOCK 7, the improvement showed a weak trend toward significance ($p=0.1$). The effect of BLOCK on error rate was not significant ($F(9,45)=0.997$, $p=0.456$).

The mean error rate was 11.3% (σ 7.2). Error rates varied from 11.7% (σ 8.8) in BLOCK 1 to 11.0% (σ 9.0) in BLOCK 10. These high error rates are attributable to a large proportion of relatively small targets in the study, that are at the limit of tracking accuracy which was not fully resolved by filtering. The larger number of *unaimed* movements in this study (40.64%) also supports this explanation.

Discussion

As in the first study, the gain functions produced by AutoGain displayed significant performance improvement form the initial condition (here up to 24%), and reasonable completion times (here around 1.5 second). Figure 12 shows the evolution of the gain functions that it produced for each participant after each block. Each function appears to converge to a particular shape, which might reflect the movement specificities and pointing strategy of each individual participant.

Finally we looked at Fitts' law regression with the data from BLOCK 10. By averaging over 10-quantiles, we obtained following equation: $MT = 0.773 + 0.211 \cdot ID$, $R^2 = 0.824$. A

previous Fitts' study of mid-air pointing using a Leap Motion device and a hand-tuned acceleration function [18] reported similar error rates (10.9% in average) and a Fitts' law regression of $MT = 0.497 + 0.483 \cdot ID$ ($R^2 = 0.94$) in a similar range of IDs (1.7 to 5.2 bits). There are a number of differences between the procedures of the two studies: in [18] participants were standing, the task was performed on a large display, participants had less training (36 trials), clicking and clutching were triggered by hand gestures with inconsistent recognition accuracy, etc. While these differences prevent formal performance comparison, we obtained overall shorter acquisition times and similar accuracy in the same range of difficulty. The fact that AutoGain could produce such a usable gain function in less than 30 minutes of use and with minimal human initialization warrants further research on automatic gain updating techniques.

PRACTICAL DEPLOYMENT

While useful in applications with dedicated pointers that can be controlled independently from the system cursor, our main objective is to allow constant, personalized gain optimization at operating system level.

Deploying AutoGain at system level requires 5 elements:

Cursor trajectory logging, which is often trivial.

Input motion logging, which is sometimes available programmatically or with dedicated hardware [23, 2]. It can also be approximated from cursor trajectory, provided sufficient knowledge about the current gain function. The same considerations apply to clutching detection: some devices provide contact-up and contact-down events. Otherwise, heuristics like duration thresholds with no input events may provide a satisfactory approximation.

Distinguishing pointing movements in real time (as opposed to e.g., idle movements or steering), for which various speed-based methods have been proposed [7, 5].

Knowledge about the intended target, which can be challenging in real use. Some approaches suggested in previous work, like crowdsourcing [7], are unsuitable for real-time optimization. Kinematic models of pointing motions such as Plamondon *e*t al.'s [25] could be used to infer the nature and characteristics (distance, width) of a pointing task using trajectories and speed profiles, although the accuracy "in the wild" of such a method remains unknown. Alternatively, accessibility APIs can detect simple actions like button presses and menu selections, as well as characteristics of the associated widget [7, 19, 5], although on a limited subset of widgets. This would cause less frequent but more trustworthy gain updates.

Altering the gain function of the system cursor. This is so far the main obstacle to applying AutoGain at system level. Most accessibility APIs can only alter the settings already available to the user. Libraries like `libpointing` [2] are not yet able to alter the gain function of an entire operating system, but greatly facilitate gain function calculations from input to output. This could be combined with low-level cursor control such as `Java Robots`[1], but this method remains to be tested.

We expect that this work, and future developments in automatic gain optimization, will encourage operating system designers to provide access to the whole gain function process.

DISCUSSION AND CONCLUSION

This paper introduces AutoGain, a novel method to obtain efficient CD gain functions automatically. It gradually adapts the gain function in order to minimize aiming error (undershooting/overshooting) at submovement level, and after every target acquisition. AutoGain requires minimal configuration (1 parameter to set manually, as opposed to 4 in previous approaches [22]), and can produce efficient gain functions in less than an hour of use.

We report two user studies. In the first one, AutoGain produced gain functions for laptop trackpads yielding performance comparable to that of widely used commercial gain functions that have been incrementally improved over decades, in about half an hour of use. In the second study, AutoGain produced gain functions for a novel input device (Leap Motion controller) for which no reference yet existed, and yielding performance comparable to previous findings where a hand-tuned gain function was used [18]. Interestingly, AutoGain produces notably distinct gain functions for different participants (see Figure 12), indicating that it has the potential to adapt to individual movements and strategies.

This proof-of-concept opens up research and application perspectives for personalized and context-dependent pointing facilitation techniques, in both existing and new systems. AutoGain could be used to accelerate the adoption of novel input devices for which no reference gain function yet exists. Pointing facilitation techniques could adapt independently, e.g. for games or drawing applications, or for different individuals sharing a computer. It might also help users with disabilities.

To improve over the current AutoGain process, we plan to further study how end-point variability changes in response to changing gain functions. This will help defining the rate at which gain updates occur (C in Equation 9), which would both accelerate the convergence of the optimization process and, possibly, the efficiency of the resulting gain functions. We also plan to apply AutoGain to more varied platforms, such as large displays controlled in mid-air, or high-precision machinery. Exploring a broad range of such systems might help us refine the initialization steps of AutoGain.

REFERENCES

1. Blanch, R., Guiard, Y., and Beaudouin-Lafon, M. Semantic pointing: Improving target acquisition with control-display ratio adaptation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, ACM (New York, NY, USA, 2004), 519–526.
2. Casiez, G., and Roussel, N. No more bricolage!: methods and tools to characterize, replicate and compare

[1] **https://docs.oracle.com/javase/8/docs/api/java/awt/Robot.html**

pointing transfer functions. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, ACM (2011), 603–614.

3. Casiez, G., Roussel, N., and Vogel, D. 1 € filter: a simple speed-based low-pass filter for noisy input in interactive systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2012), 2527–2530.
4. Casiez, G., Vogel, D., Balakrishnan, R., and Cockburn, A. The impact of control-display gain on user performance in pointing tasks. *HUMANCOMPUTER INTERACTION 23*, 3 (2008), 215–250.
5. Chapuis, O., Blanch, R., and Beaudouin-Lafon, M. Fitts' Law in the Wild: A Field Study of Aimed Movements. Tech. rep., Dec. 2007. LRI Technical Repport Number 1480, Univ. Paris-Sud, 11 pages.
6. Crossman, E., and Goodeve, P. Feedback control of hand-movement and fitts' law. *The Quarterly Journal of Experimental Psychology 35*, 2 (1983), 251–278.
7. Evans, A., and Wobbrock, J. Taming wild behavior: The input observer for obtaining text entry and mouse pointing measures from everyday computer use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, ACM (New York, NY, USA, 2012), 1947–1956.
8. Fradet, L., Lee, G., and Dounskaia, N. Origins of submovements during pointing movements. *Acta psychologica 129*, 1 (2008), 91–100.
9. Frees, S., Kessler, G. D., and Kay, E. Prism interaction for enhancing control in immersive virtual environments. *ACM Trans. Comput.-Hum. Interact. 14*, 1 (May 2007).
10. Gallo, L., Ciampi, M., and Minutolo, A. Smoothed pointing: A user-friendly technique for precision enhanced remote pointing. In *Proceedings of the 2010 International Conference on Complex, Intelligent and Software Intensive Systems*, CISIS '10, IEEE Computer Society (Washington, DC, USA, 2010), 712–717.
11. Haque, F., Nancel, M., and Vogel, D. Myopoint: Pointing and clicking using forearm mounted electromyography and inertial motion sensors. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, ACM (New York, NY, USA, 2015), 3653–3656.
12. Hart, S. G. Nasa-task load index (nasa-tlx); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, vol. 50, Sage Publications (2006), 904–908.
13. Kalman, R. E. A new approach to linear filtering and prediction problems. *Journal of basic Engineering 82*, 1 (1960), 35–45.
14. König, W. A., Gerken, J., Dierdorf, S., and Reiterer, H. Adaptive pointing: Implicit gain adaptation for absolute pointing devices. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, ACM (New York, NY, USA, 2009), 4171–4176.
15. Kozlov, Y., and Weinkauf, T. Persistence1d. Code, November 2013. Retrieved June 2, 2016 from **`https://people.mpi-inf.mpg.de/~weinkauf/notes/persistence1d.html`**.
16. Lee, B., and Bang, H. A kinematic analysis of directional effects on mouse control. *Ergonomics 56*, 11 (2013), 1754–1765.
17. Lee, B., and Bang, H. A mouse with two optical sensors that eliminates coordinate disturbance during skilled strokes. *Human–Computer Interaction 30*, 2 (2015), 122–155.
18. Liu, M., Nancel, M., and Vogel, D. Gunslinger: Subtle arms-down mid-air interaction. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15, ACM (New York, NY, USA, 2015), 63–71.
19. Malacria, S., Scarr, J., Cockburn, A., Gutwin, C., and Grossman, T. Skillometers: Reflective widgets that motivate and help users to improve performance. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology*, UIST '13, ACM (New York, NY, USA, 2013), 321–330.
20. Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E., and Keith Smith, J. Optimality in human motor performance: ideal control of rapid aimed movements. *Psychological review 95*, 3 (1988), 340.
21. Nancel, M., Chapuis, O., Pietriga, E., Yang, X.-D., Irani, P. P., and Beaudouin-Lafon, M. High-precision pointing on large wall displays using small handheld devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, ACM (New York, NY, USA, 2013), 831–840.
22. Nancel, M., Pietriga, E., Chapuis, O., and Beaudouin-Lafon, M. Mid-air pointing on ultra-walls. *ACM Trans. Comput.-Hum. Interact. 22*, 5 (Aug. 2015), 21:1–21:62.
23. Nancel, M., Vogel, D., and Lank, E. Clutching is not (necessarily) the enemy. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, ACM (New York, NY, USA, 2015), 4199–4202.
24. Phillips, J. G., Triggs, T. J., and Meehan, J. W. Forward/up directional incompatibilities during cursor placement within graphical user interfaces. *Ergonomics 48*, 6 (2005), 722–735.
25. Plamondon, R. The kinematic theory: a new window to study and analyze simple and complex human movements. *Behavioral and Brain Sciences 20*, 02 (1997), 325–343.
26. Weichert, F., Bachmann, D., Rudak, B., and Fisseler, D. Analysis of the accuracy and robustness of the leap motion controller. *Sensors 13*, 5 (2013), 6380–6393.

27. Wisleder, D., and Dounskaia, N. The role of different submovement types during pointing to a target. *Experimental Brain Research 176*, 1 (2007), 132–149.

28. Wobbrock, J. O., Cutrell, E., Harada, S., and MacKenzie, I. S. An error model for pointing based on fitts' law. In *Proceedings of the SIGCHI conference on human factors in computing systems*, ACM (2008), 1613–1622.

29. Wobbrock, J. O., Fogarty, J., Liu, S.-Y. S., Kimuro, S., and Harada, S. The angle mouse: Target-agnostic dynamic gain adjustment based on angular deviation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, ACM (New York, NY, USA, 2009), 1401–1410.

30. Yun, J., Lim, Y., Kim, K., and Song, S. Interactivity crafter: An interactive input-output transfer. *Archives of Design Research 28*, 3 (2015), 21–37.

31. Zhai, S., Kong, J., and Ren, X. Speed–accuracy tradeoff in fitts law taskson the equivalency of actual and nominal pointing precision. *International journal of human-computer studies 61*, 6 (2004), 823–856.